

INTRODUCING INNER NESTED SAMPLING

H.R.N. VAN ERP, R.O. LINGER, AND P.H.A.J.M. VAN GELDER

ABSTRACT. In this paper we will give a Monte Carlo algorithm by which the moments of a functions of Dirichlet probability distributions can be estimated. This algorithm is called Inner Nested Sampling and is an implementation of Skilling's general Nested Sampling framework.

1. INTRODUCTION

This paper is for the benefit of those readers who are interested in finding moments of functions of Dirichlet probability distributions. In a data analysis we will, typically, have some set of count data, say,

$$D = \{r_1, r_2, \dots, r_n\},$$

which by way of a Dirichlet distribution,

$$p(\boldsymbol{\theta}) = \frac{(\sum r_i - 1)!}{(r_1 - 1)!(r_1 - 2)! \dots (r_n - 1)!} \theta_1^{r_1-1} \theta_2^{r_2-1} \dots \theta_n^{r_n-1}, \quad (1.1)$$

points to some set of underlying probabilities $\boldsymbol{\theta}$, where

$$\boldsymbol{\theta} = \{\theta_1, \theta_2, \dots, \theta_n\}.$$

Stated differently, the $\boldsymbol{\theta}$ are not directly observable, they can only be inferred. So, if we wish to assign a function u to $\boldsymbol{\theta}$, then we have to take our uncertainty, in regards to the actual value of the vector $\boldsymbol{\theta}$, into account. But if we do so, then this will give us highly dimensional and highly intractable integrals. In order to evaluate these integrals, we will invoke the Nested Sampling *framework*.

In order to *implement* this Nested Sampling framework, we will then proceed to introduce an algorithm, called Inner Nested Sampling. This algorithm is, for the specific problem of transforming the probability distribution $\boldsymbol{\theta}$ to a probability distribution of the function $u(\boldsymbol{\theta})$, the optimal implementation, in that it takes into the specific geometry of the probability distribution (1.1).

2. WHY NESTED SAMPLING?

In order to answer question why one ought to use Nested Sampling (NS) as their Monte Carlo framework of choice, we quote Skilling himself (personal communication, 2013):

NS is *designed* to get the evidence value, which is equation #1 of inference.

$$(1) \quad Z = \text{INTEGRAL } d\theta \text{ JointProb}(\theta, \text{data})$$

$$(2) \quad \text{Posterior}(\theta) = \text{JointProb}(\theta, \text{data}) / Z$$

Logically, Z precedes the posterior and should never be ignored. It's the more important half of Bayes, not only in the algebra, but also because there's no point in bothering to generate a posterior if the evidence falls much short of some other plausible model.

Stated differently, let $p(\theta, D)$ be the product of some prior $p(\theta)$ of some unknown set of parameters θ and some likelihood function $L(\theta, D)$ of some data set D :

$$p(\theta, D) = p(\theta) L(\theta, D). \quad (2.1)$$

Then the integral

$$Z = \int p(\theta, D) d\theta \quad (2.2)$$

is the *evidence* measure which may be used to differentiate between competing models, by way of Bayesian model selection. Furthermore, the posterior of the parameters θ , given the data D , that is, $p(\theta|D)$, is given as:

$$p(\theta|D) = \frac{p(\theta, D)}{Z}. \quad (2.3)$$

The Nested Sampling algorithm [1, 2] is specifically designed to evaluate the integral (2.2), giving us an estimate of the evidence Z . Furthermore, it also provides us with a set of representative samples from the posterior (2.3), which may function as a proxy for that posterior.

For those cases where the integral (2.2) may be evaluated analytically, one will have no need for the Nested Sampling algorithm. However, for those problems where the integral (2.2) is both intractable and highly dimensional, there one will have to take his recourse to Nested Sampling, in order to be able to evaluate the evidence (2.2) and obtain a set of representative samples from the desired posterior (2.3).

The larger the evidence, the larger the adequacy of (2.1), in terms of both parsimony and data fit. Stated differently, if we are considering some $p(\theta, D)$, then we only need to bother to construct a posterior for the model which has the largest evidence value. For what use are our (admittedly accurate) parameter estimates, as captured in the posterior distribution, if the model that leads us to these estimates is woefully inadequate?

Before proceeding to the Nested Sampling algorithm, we will elaborate on the evidence measure and its role in Bayesian model selection. This is done for the benefit of those of the readers who are not yet well acquainted with these concepts.

3. THE EVIDENCE IN BAYESIAN MODEL SELECTION

Bayesian statistics has four fundamental constructs, namely, the prior, the likelihood, the posterior, and the evidence. These constructs are related in the following way:

$$\text{posterior} = \frac{\text{prior} \times \text{likelihood}}{\text{evidence}}. \quad (3.1)$$

As an aside, any student of Bayesian probability theory will tend to have a firm grip on the concepts of a prior, likelihood, and posterior. However, the concept of evidence is less universally known. A possible explanation for this is that most people (used to) come to Bayesian probability theory by way of the more compact relationship

$$\text{posterior} \propto \text{prior} \times \text{likelihood}, \quad (3.2)$$

which does not make any explicit mention of the evidence construct¹.

In what follows, we will employ in our analysis the correct, though notationally more cumbersome, relation (3.1), and forgo of the more compact, but incomplete, Bayesian shorthand (3.2). This is done so the reader may develop some feeling for the evidence construct, and how this construct relates to the other three Bayesian constructs of prior, likelihood, and posterior.

Let $p(\theta|I)$ be the prior of some parameter θ , where I is the prior information regarding the unknown θ which we have to our disposal. Let $p(D|\theta, M)$ be the probability of the data D conditional on the value of parameter θ and some likelihood model M which is used²; the probability of the data is also known as the likelihood of the parameter θ . Let $p(\theta|D, M, I)$ be the posterior distribution of the parameter θ , conditional on the data D , the likelihood model M , and the prior information model I . Then

$$p(\theta|D, M, I) = \frac{p(\theta|I)p(D|\theta, M)}{\int p(\theta|I)p(D|\theta, M)d\theta} = \frac{p(\theta|I)p(D|\theta, M)}{p(D|M, I)}, \quad (3.3)$$

where

$$p(D|M, I) = \int p(\theta, D|M, I)d\theta = \int p(\theta|I)p(D|\theta, M)d\theta \quad (3.4)$$

is the evidence, that is, the marginalized likelihood, of both the likelihood model M and the prior information model I . We now will show how this evidence may be used in Bayesian model selection.

If we have a set of likelihood models M_j we wish to choose from, and just the one prior information model I , then we may do so by computing the evidence values $p(D|M_j, I)$.

¹See for example [4] throughout.

²Note that M is a label that points to some likelihood model, and not so much a parameter whose value we wish to determine.

Let $p(M_j)$ and $p(M_j|D, I)$ be, respectively, the prior and posterior probability of the likelihood model M_j . Then the posterior probability distribution of these likelihood models is given as

$$p(M_j|D, I) = \frac{p(M_j)p(D|M_j, I)}{\sum_j p(M_j) \cdot p(D|M_j, I)} \quad (3.5)$$

Note that if $p(M_j) = p(M_k)$ for all j and k , then we have that (3.5) reduces to

$$p(M_j|D, I) = \frac{p(D|M_j, I)}{\sum_j p(D|M_j, I)}. \quad (3.6)$$

Stated differently, if we assign equal prior probabilities to our different models, then these models may be ranked by their respective evidence values $p(D|M_j, I)$, [5].

4. NESTED SAMPLING: THE IDEA

By reducing any k -variate function f to a corresponding monotonic descending univariate function g , and by using order statistics, the integral of any k -variate function f may be evaluated using a Monte Carlo sampling scheme called Nested Sampling, [2].

Say, we wish to numerically evaluate a bivariate distribution $f(x, y)$, where

$$f(x, y) = \frac{\sqrt{1 - 0.7^2}}{2\pi} \exp \left[-\frac{1}{2} (x^2 + 1.4xy + y^2) \right] \quad (4.1)$$

where $-5 \leq x, y \leq 5$, Figure 1. The total volume under the curve (4.1) is given by




FIGURE 1. Plot of Function f

the integral

$$\int_{-5}^5 \int_{-5}^5 \frac{\sqrt{1 - 0.7^2}}{2\pi} \exp \left[-\frac{1}{2} (x^2 + 1.4xy + y^2) \right] dx dy = 0.9993 \quad (4.2)$$

We may evaluate the integral (4.1) through brute computational force. Say, we partition the x, y -plane in little squares with areas $dx_j dy_k = 0.25$, for $j = 1, \dots, 20$ and $k = 1, \dots, 20$. Then define the center of these areas as $(\tilde{x}_j, \tilde{y}_k)$, and compute the strips of volume V_{jk} as

$$V_{jk} = f(\tilde{x}_j, \tilde{y}_k) dx_j dy_k \quad (4.3)$$

In Figure 2 we give all the volume elements V_{jk} together.




FIGURE 2. Volume Elements of Function f

Summating the volume elements (4.3), the total volume under the curve (4.1) may be approximated as

$$\text{volume} \approx \sum_{j=1}^{20} \sum_{k=1}^{20} V_{jk} = 0.9994 \quad (4.4)$$

The 3-dimensional volume elements V_{jk} may be mapped to corresponding 2-dimensional area elements A_i , without any loss of generality. In order to demonstrate this, we introduce the following notation

$$dw_i = dx_j dy_k, \quad f_i = f(\tilde{x}_j, \tilde{y}_k) \quad (4.5)$$


where index i is a function of the indices j and k :

$$i \equiv (j - 1) 20 + k \quad (4.6)$$

and $i = 1, \dots, 400$. Using (4.5), we may rewrite (4.3) as

$$V_{jk} = f(\tilde{x}_j, \tilde{y}_k) dx_j dy_k = f_i dw_i = A_i \quad (4.7)$$


In Figure 3 we give all the 400 area elements A_i together.

FIGURE 3. Area Elements of Function f

Since (4.7) is equivalent to (4.3), we have that the mapping of the 3-dimensional volume elements, V_{jk} , to their corresponding 2-dimensional area elements, A_i , has not led to any loss of information; that is,


$$\text{area} = \sum_{i=1}^{400} A_i = \sum_{j=1}^{20} \sum_{k=1}^{20} V_{jk} = \text{volume} \quad (4.8)$$

We may, trivially, rearrange the elements A_i in Figure 3 in descending order, Figure 4.

FIGURE 4. Ordered Area Elements of Function f

All these rectangular elements have a base of $dw = dx \, dy = 0.25$. So, being that there are 400 area elements, we might view Figure 4 as a representation of some hypothetical monotonic descending function $g(w)$, where $0 \leq w \leq 100$, Figure 5.

What we have accomplished is that we have mapped 3-dimensional volume elements, Figure 2, to 2-dimensional area elements, Figure 3, after which we have rearranged the area elements, Figure 4, so as to get a monotonic descending function

FIGURE 5. Plot of Hypothetical Function g

$g(w)$, Figure 5. If we integrate the univariate function $g(w)$, Figure 5, we get the volume we are looking for.

Any k -variate function may be reduced to its corresponding monotonic descending univariate function $g(w)$. We will see that the procedure of Nested Sampling is based upon the equivalence between any k -variate function and its corresponding univariate representation $g(w)$.

4.1. Estimating abscissa's. Skilling's Nested Sampling framework is unique in that it is a Monte Carlo scheme that uses probability theory to its advantage, [2].

If we have a value of $g(w)$, without knowing the value of w . Then the only thing we know about w is that it must lie somewhere in the region $0 \leq w \leq W$, where W is the scalar value of the volume on which the k -variate function is defined.

So, if, based on this information alone, we were asked to assign a probability distribution to the actual value of the abscissa w that corresponds with the ordinate $g(w)$, then we would have to assign a univariately uniformly distribution on the domain $[0, W]$, that is,

$$p(w) = \frac{1}{W}, \quad 0 \leq w \leq W. \quad (4.9)$$

Consequently, the unknown abscissa value w has an expectation value of

$$E(w) = \frac{W}{2} \quad (4.10)$$

and a standard deviation of

$$\text{std}(w) = \frac{W}{2\sqrt{3}}. \quad (4.11)$$

Let $\mathbf{x} = (x_1, \dots, x_k)$ be a point in the k -dimensional domain. If we have N random samples of the k -variate domain, that is, $\{\mathbf{x}_1, \dots, \mathbf{x}_N\}$. Then, by evaluating these points, we have a set of N random ordinates, that is, $\{f(\mathbf{x}_1), \dots, f(\mathbf{x}_N)\}$. Because of the equivalence between any k -variate function f and its univariate

representation g , we may write

$$g(w_n) = f(\mathbf{x}_n), \quad n = 1, \dots, N \quad (4.12)$$

So, by construction, the set of N random ordinates $\{f(\mathbf{x}_1), \dots, f(\mathbf{x}_N)\}$ corresponds with the set of random ordinates $\{g(w_1), \dots, g(w_N)\}$.

Even though we have the ordinate $g(w_n)$, by way of (4.12), we do not know the abscissa w_n that goes with this ordinate. As we sample the k -variate domain \mathbf{x} , rather than the univariate domain w . However, the one thing we do know is that the smallest realisation of $g(w_n)$ corresponds with the greatest value of w_n . This is because function g is a monotonic descending function, see Figure 5. It follows that we may use an order distribution for the unknown value $\max(w_n)$, or, for short, w_{\max} .

The order probability distribution for the largest draw, w_{\max} , in a sample of size N from the uniform distribution (4.12) is

$$p(w_{\max}) = N \left(\frac{w_{\max}}{W} \right)^{N-1} \frac{1}{W} \quad (4.13)$$

where both N and W are known. The order distribution (4.13) has a mean of

$$E(w_{\max}) = \left(1 - \frac{1}{N+1} \right) W \quad (4.14)$$

and a standard deviation of

$$\text{std}(w_{\max}) = \sqrt{\frac{N}{(N+1)^2 (N+2)}} W \quad (4.15)$$

We see in (4.15) that the standard deviation, that is, our uncertainty regarding the unknown value of w_{\max} falls off with a factor N . We will see that (4.14) and (4.15) form the backbone of the Nested Sampling framework.

Note that Skilling himself does not arrive at Nested Sampling algorithm by way of order statistics, but by way of the closely related probability distribution of the shrinkage ratio of the admissible domain of the unknown univariate representation g of the multivariate function f :

$$t = \frac{w_{\max}}{W}. \quad (4.16)$$

This shrinkage ratio has a probability distribution

$$p(t) = N t^{N-1}. \quad (4.17)$$

And we may compare (4.13) with (4.17), to see that both probability distributions are equal, up to a change of variable, which scales the maximum abscissa value w_{\max} to the fraction t , (4.16).

4.2. Nested Sampling: The Algorithm. In this initial version of the Nested Sampling framework we will not protect against under- and overflow. We will just focus here on the basic philosophy which underlies Nested Sampling.

In what follows we reserve subscripts for sample membership and superscripts for the enumeration of the algorithmic step.

Step 1

Find N random values \mathbf{x} in the k -variate domain and evaluate, that is, find

$$f(\mathbf{x}_n^{(1)}), \quad n = 1, \dots, N \quad (4.18)$$

Since we may perform the steps as shown in Figure 1 through Figure 5, it holds, by way of construction, that the N values of $f(\mathbf{x}_n^{(1)})$ correspond with N values of $g(w_n)$, (4.12).

As stated before, we cannot explicitly map the abscissa \mathbf{x} to the corresponding abscissa of the monotonic descending function, w . We could link the volume elements (4.3) to the area elements (4.10) using brute computational force. But this would amount to regular numerical integration and, thus, we would still be faced with the curse of dimensionality. So, instead, we use (4.14) to statistically approximate the w which corresponds with the smallest sampled value of $f(\mathbf{x}_n^{(1)})$ and, thus, get our first coordinate $[w^{(1)}, g^{(1)}]$ of the unknown function g , Figure 5. The abscissa and ordinate of this coordinate are, respectively,

$$w^{(1)} = \left(1 - \frac{1}{N+1}\right) W \quad (4.19)$$

and

$$g^{(1)} = \min f(\mathbf{x}_n^{(1)}) \quad (4.20)$$

As the error of the statistically approximated $w^{(1)}$ falls off with a factor N , (4.15), the accuracy of the approximation (4.19) will increase as N , the number of random samples, increases.

Finally, we approximate the integral to the right of $w^{(1)}$ of our constructed function g , Figure 5, as

$$A^{(1)} = (W - w^{(1)}) g^{(1)} = \frac{W}{N+1} g^{(1)} \approx \int_{w_1}^W g(w) dw \quad (4.21)$$

and set the evaluated integral at the first step, $Z^{(1)}$, to

$$Z^{(1)} = A^{(1)} \quad (4.22)$$

Step 2

We again find N random values \mathbf{x} in the k -variate domain and evaluate. But now we constrain the random values \mathbf{x} in that their mapping to the function f be equal or greater than the minimum value of the mapping of the previous mapping, that is, we sample the $\mathbf{x}_n^{(2)}$ under the constraint

$$f(\mathbf{x}_n^{(2)}) \geq \min f(\mathbf{x}_m^{(1)}), \quad n, m = 1, \dots, N \quad (4.23)$$

In the first iteration our state of knowledge was that $0 \leq w \leq W$. In the second iteration we sample $\mathbf{x}_n^{(2)}$ under constraint (4.23), which is equivalent to sampling w under the constraint, (4.20),

$$g(w) \geq g^{(1)} \quad (4.24)$$

where g is a monotonic descending function in w . So, the updated state of knowledge is that $0 \leq w \leq w^{(1)}$.

Using (4.14) again, but replacing W , the initial upperbound of w , with $w^{(1)}$, the new upperbound of w , we may approximate the second coordinate $[w^{(2)}, g^{(2)}]$ of the unknown function g :

$$w^{(2)} = \left(1 - \frac{1}{N+1}\right) w^{(1)} = \left(1 - \frac{1}{N+1}\right)^2 W \quad (4.25)$$

and

$$g^{(2)} = \min f(\mathbf{x}_n^{(2)}) \quad (4.26)$$

We approximate the area of the integral between $w^{(2)}$ and $w^{(1)}$, where $w^{(2)} < w^{(1)}$, as

$$A^{(2)} = (w^{(1)} - w^{(2)}) g^{(2)} = \frac{w^{(1)}}{N+1} g^{(2)} \approx \int_{w^{(2)}}^{w^{(1)}} g(w) dw \quad (4.27)$$

and set the evaluated integral at the second step, $Z^{(2)}$, to

$$Z^{(2)} = A^{(1)} + A^{(2)} \quad (4.28)$$

Step t

In iteration step t we approximate coordinate $[w^{(t)}, g^{(t)}]$ of the unknown function g by way of

$$w^{(t)} = \left(1 - \frac{1}{N+1}\right) w^{(t-1)} = \left(1 - \frac{1}{N+1}\right)^t W \quad (4.29)$$

and

$$g^{(t)} = \min f(\mathbf{x}_n^{(t)}) \quad (4.30)$$

where the $\mathbf{x}_n^{(t)}$ have been sampled under the constraint

$$f(\mathbf{x}_n^{(t)}) \geq \min f(\mathbf{x}_m^{(t-1)}), \quad n, m = 1, \dots, N \quad (4.31)$$

The area of the integral between $w^{(t)}$ and $w^{(t-1)}$, where $w^{(t)} < w^{(t-1)}$, is approximated as

$$A^{(t)} = \left(w^{(t-1)} - w^{(t)}\right) g^{(t)} = \frac{w^{(t-1)}}{N+1} g^{(t)} \approx \int_{w^{(t)}}^{w^{(t-1)}} g(w) dw \quad (4.32)$$

The evaluated integral at the step t , $Z^{(t)}$, is updated as

$$Z^{(t)} = Z^{(t-1)} + A^{(t)} = \sum_{i=1}^t A^{(i)} \quad (4.33)$$

Termination step

Because of the identity (4.29)

$$w^{(t)} = \left(1 - \frac{1}{N+1}\right)^t W$$

we have that $w^{(t)} \rightarrow 0$ as $t \rightarrow \infty$. If f is bounded, then by construction so is g . It then follows from (4.32) that bounded functions f or, equivalently, bounded g

$$A^{(t+1)} = \frac{w^{(t)}}{N+1} g^{(t+1)} \rightarrow 0 \quad (4.34)$$

as $t \rightarrow \infty$. The maximum of f , the function whose integral we wish to evaluate, is also the bound of g , its univariate representation. This implies the following inequality

$$g^{(t+1)} \leq \max f(\mathbf{x}) \quad (4.35)$$

By way of (4.32) and (4.35), we then have that

$$A^{(t+1)} \leq \frac{w^{(t)}}{N+1} \max f(\mathbf{x}) \quad (4.36)$$

So, if we can determine the maximum of f , we can use as a possible terminating condition the point where the upperbound of $A^{(t+1)}$ does not contribute more than N^2 part to $Z^{(t+1)}$

$$\frac{w_t}{N+1} \max f(\mathbf{x}) < \frac{Z^{(t+1)}}{N^2} \quad (4.37)$$

Alternatively, if we are unable determine the maximum of f , then we take as a stopping criterium the point where $A^{(t+1)}$ itself does not contribute more than N^2 part to $Z^{(t+1)}$

$$A^{(t+1)} < \frac{Z^{(t+1)}}{N^2} \quad (4.38)$$

However, (4.38) comes with the caveat that the function f , or, equivalently, g , might still have some unexplored regions having large enough values $w^{(q+1)}$, where $q > t$, to take

$$A^{(q+1)} = \frac{w^{(q)}}{N+1} g^{(q+1)} > \frac{Z^{(q+1)}}{N^2} \quad (4.39)$$

which would imply a premature termination of the algorithm.

4.3. Optimal implementation. If one is to implement the steps in the previous section one would get an algorithm very differently from the pseudo-code given by Skilling, [2]. These differences are mainly differences in implementation.

The reason for us to first give the naive algorithm was to point to reader to the elegant and simple idea behind Nested Sampling, without losing ourselves too much in the technicalities of optimal implementations.

However, with the core idea behind Nested Sampling demonstrated, we now will treat in the next two paragraphs the points of optimal implementation.

4.3.1. Reduction of computational cost. In the previous treatment of the Nested Sampling framework, N new samples of \mathbf{x} were drawn at each iteration step t , under the constraint

$$f(\mathbf{x}_n^{(t)}) \geq \min f(\mathbf{x}_m^{(t-1)}), \quad n, m = 1, \dots, N \quad (4.40)$$

Now, instead of drawing N new samples at each sampling step t , we may also realize that in iteration $(t-1)$, we already had $(N-1)$ samples of f at our disposal that satisfied the constraint (4.40). Namely, those samples in iteration step $(t-1)$ smaller than

$$\min f(\mathbf{x}_m^{(t-1)}).$$

If we keep these $(N-1)$ samples, then we only need to sample one additional ordinate value which satisfies the constraint (4.40), in order to obtain our needed sample of N objects.

So, after each iteration t we discard one object from our sample of N objects. This discarded object becomes $g^{(t)}$. The $(N-1)$ surviving objects are taken to the next iteration and an additional object is sampled under constraint

$$f(\mathbf{x}_n^{(t+1)}) \leq g^{(t)}.$$

This implementation reduces the computational costs of Nested Sampling with an order of magnitude of N .

4.3.2. Guarding against under- and overflow. In many problems $Z^{(t)}$ may become so large that computational overflow may occur, that is, there is no longer a floating number representation possible for its value. To remedy this problem we will have to work with the $\log Z^{(t)}$. Furthermore, We also have that certain functions f have certain values in their domain so small that computational underflow may occur. So, in what follows we will evaluate $\log f$ instead of f . Likewise, because of the fact that $w_t \rightarrow 0$ as $t \rightarrow \infty$, we have that for sufficiently large t computational underflow may occur, that is, w_t may become so small that there is no longer a

floating number representation possible for its value. To remedy the latter situation we will go from the w scale to the $\log w$ scale.

To go to the $u = \log w$ scale, we will have to make a proper change of variable for the order distribution (4.13). We have that

$$du = \frac{dw}{w}, \quad w = \exp u \quad (4.41)$$

Substituting (4.41) into (4.13), we may obtain

$$p(u_{\max}) = \frac{N}{W^N} \exp(Nu_{\max}) \quad (4.42)$$

for $-\infty < u_{\max} \leq \log W$. with mean standard deviation of

$$E(u_{\max}) = -\frac{1}{N} + \log W \quad (4.43)$$

and

$$\text{std}(u_{\max}) = \frac{1}{N} \quad (4.44)$$

With a repeated application of (4.43) we may find the limits of the u scale after the t th iteration to be

$$u^{(t)} = -\frac{t}{N} + \log W \quad (4.45)$$

From (4.45), it follows that that the width of the t th interval on the original w scale may be written as

$$\begin{aligned} w^{(t-1)} - w^{(t)} &= \exp(u^{(t-1)}) - \exp(u^{(t)}) \\ &= \exp\left(-\frac{t-1}{N} + \log W\right) - \exp\left(-\frac{t}{N} + \log W\right) \\ &= W \left[\exp\left(-\frac{t-1}{N}\right) - \exp\left(-\frac{t}{N}\right) \right] \\ &= W \exp\left(-\frac{t-1}{N}\right) \left[1 - \exp\left(-\frac{1}{N}\right) \right] \end{aligned} \quad (4.46)$$

Consequently, we have

$$\log(w^{(t-1)} - w^{(t)}) = -\frac{t-1}{N} + \log W + \log \left[1 - \exp\left(-\frac{1}{N}\right) \right] \quad (4.47)$$

We then have that the log of the area element at the t th iteration is the sum

$$\begin{aligned} \log A^{(t)} &= \log \left(w^{(t-1)} - w^{(t)} \right) + \log g^{(t)} \\ &= -\frac{t-1}{N} + \log W + \log \left[1 - \exp \left(-\frac{1}{N} \right) \right] + \log g^{(t)} \end{aligned} \quad (4.48)$$

With (4.48) the computation of $A^{(t)}$ is sufficiently protected from underflow of both $w^{(t)}$ and f , or, equivalently, g .

We may update $\log Z^{(t)}$ with $\log A^{(t)}$ using the formula for logarithmic addition, (Skilling, 2006),

$$\begin{aligned} \log (\exp x + \exp y) &= \log \{ \exp x [1 + \exp (y - x)] \} \\ &= x + \log [1 + \exp (y - x)] \end{aligned} \quad (4.49)$$

If we set $x = \log Z^{(t-1)}$ and $y = \log A^{(t)}$, then (4.49) gives us

$$\log Z^{(t)} = \log Z^{(t-1)} + \log \left[1 + \exp \left(\log A^{(t)} - \log Z^{(t-1)} \right) \right] \quad (4.50)$$

With (4.50) the computation of $\log Z^{(t)}$ is sufficiently protected from the overflow of $\log Z^{(t-1)}$.

To summarize, in order to protect the algorithm from under- and overflow, which may easily occur in actual problems, the original algorithm remains unchanged, except that each iteration $\log A^{(t)}$ is computed, instead of $A^{(t)}$, by way of (4.48), and $\log Z^{(t)}$ is updated, instead of $Z^{(t)}$, by (4.50). Note that the termination conditions (4.37) and (4.38) transform, respectively, to

$$-\frac{t}{N} + \log W + \log \left[1 - \exp \left(-\frac{1}{N} \right) \right] + \log [\max f(\mathbf{x})] < \log Z^{(t+1)} - 2 \log N \quad (4.51)$$

if we know the maximum of the function f , and

$$\log A^{(t)} < \log Z^{(t+1)} - 2 \log N \quad (4.52)$$

if we do not.

5. APPLYING NESTED SAMPLING TO UNSCALED POSTERiors

Nested Sampling is a general purpose algorithm for the evaluation of the integrals of multivariate functions. A special class of multivariate functions are unscaled posterior distributions.

Let $p(x, y)$ be some prior distribution of the unknown parameters x and y . Let $p(D|x, y)$ be the probability of some observed data set D given the parameters x and y , or, equivalently, the likelihood $L(x, y)$. Then the unscaled posterior distribution

is

$$p(x, y, D) = p(x, y) p(D | x, y) = p(x, y) L(x, y) \quad (5.1)$$

If we integrate $p(x, y, D)$ over the domain of x and y we get the marginal probability of the observed data set D

$$p(D) = \int \int p(x, y, D) dx dy = \int \int p(x, y) p(D | x, y) dx dy \quad (5.2)$$

The marginal probability $p(D)$ is also called the *evidence*. Scaling $p(x, y, D)$ with $p(D)$ we get, by way of the product rule, the scaled posterior distribution:

$$p(x, y | D) = \frac{p(x, y, D)}{p(D)} \quad (5.3)$$

Furthermore, the evidence is an important quantity in its own right. It may be used for model selection. Say, we have two competing likelihood models, L_1 and L_2 . Combining (5.1) and (5.2), we may compare the respective evidences of the competing likelihood models:

$$p_1(D) = \int \int p(x, y) L_1(x, y) dx dy \quad (5.4)$$

$$p_2(D) = \int \int p(x, y) L_2(x, y) dx dy$$

Then the likelihood model that fits the data best will be the model that gives us the greatest marginal probability of obtaining the observed data set D . The evaluation of the integral (5.1) is a non-trivial matter, especially for highly variate likelihood functions. It was with this application in mind that Skilling developed his Nested Sampling framework.

We now proceed to apply the Nested Sampling framework to (5.1). Let the function f be

$$f(x, y) = p(x, y) L(x, y) \quad (5.5)$$

Let the prior $p(x, y)$ be the uniform distribution over the domain of x and y . If this domain has a total area of W , then we have that

$$p(x, y) = \frac{1}{W} \quad (5.6)$$

Substituting (5.6) into (5.5), and substituting the resulting function f into (4.48), we get

$$\begin{aligned}
\log A^{(t)} &= -\frac{t-1}{N} + \log W + \log \left[1 - \exp \left(-\frac{1}{N} \right) \right] + \log g^{(t)} \\
&= -\frac{t-1}{N} + \log W + \log \left[1 - \exp \left(-\frac{1}{N} \right) \right] - \log W + \log \left(\min L_n^{(t)} \right) \\
&= -\frac{t-1}{N} + \log \left[1 - \exp \left(-\frac{1}{N} \right) \right] + \log \left(\min L_n^{(t)} \right)
\end{aligned} \tag{5.7}$$

where $\min L_n^{(t)}$ is the smallest likelihood from the sample of N likelihoods at the t th step, which were sampled under the constraint that

$$L_n^{(t)} \leq \min L_m^{(t-1)} \tag{5.8}$$

The Nested Sampling framework is most often presented in the form of (5.7), which is the special case of f being an unscaled posterior which takes as its prior the uniform distribution of the parameter space. However, the framework is much more general in that it is applicable to any function f , not just unscaled posterior distributions.

Last but not least, if f is an unscaled posterior, then the Nested Sampling framework not only evaluates the evidence. It also gives us a Monte Carlo proxy of the posterior of interest. Let T be the termination step of the Nested Sampling run. Let $\mathbf{x}^{(t)}$ be the points in the k -variate domain that corresponded with the $\mathbf{x}_n^{(t)}$ for which

$$g^{(t)} = \min p(\mathbf{x}_n^{(t)}) L(\mathbf{x}_n^{(t)}) \tag{5.9}$$

where $t = 1, \dots, T$ and $n = 1, \dots, N$. Then we may assign to these $\mathbf{x}^{(t)}$ the probability weights, (4.32) and (4.33),

$$P(\mathbf{x}^{(t)}) = \frac{A^{(t)}}{Z^{(T)}} \tag{5.10}$$

This leaves us with a set of weighted random Monte Carlo samples. The weighted samples are a proxy for the posterior distribution of interest, [2].

Let the posterior $p(\mathbf{x} | D)$ be given as

$$p(\mathbf{x} | D) = \frac{p(\mathbf{x}) L(\mathbf{x})}{\int p(\mathbf{x}) L(\mathbf{x}) d\mathbf{x}} \tag{5.11}$$

Let h be some function defined on the k -variate domain of the \mathbf{x} . Then we may approximate the weighted h by way of Nested Sampling as, (5.10),

$$\int h(\mathbf{x}) p(\mathbf{x}|D) d\mathbf{x} \approx \sum_{t=1}^T h(\mathbf{x}^{(t)}) P(\mathbf{x}^{(t)}) = \sum_{t=1}^T h(\mathbf{x}^{(t)}) \frac{A^{(t)}}{Z^{(T)}} \quad (5.12)$$

It is (5.12) that will enable us to compute the first and second moments of the relevances defined on a Dirichlet distribution.

6. THE ISSUE OF CONFIDENCE BOUNDS

Say we wish to assign the function u to a bivariate probability distributions of the form:

	b_1	b_2	
a_1	θ_1	θ_3	$\theta_1 + \theta_3$
a_2	θ_2	θ_4	$\theta_2 + \theta_4$
	$\theta_1 + \theta_2$	$\theta_3 + \theta_4$	$\theta_1 + \theta_2 + \theta_3 + \theta_4 = 1$

TABLE 1. Distribution 1

Then, in practice, we only have indirectly access, by way of our data, to the probability distribution of interest; that is, rather than the known probabilities, Table 1, we only have some observed count data, Table 2:

	b_1	b_2	
a_1	r_1	r_3	$r_1 + r_3$
a_2	r_2	r_4	$r_2 + r_4$
	$r_1 + r_2$	$r_3 + r_4$	$r_1 + r_2 + r_3 + r_4 = n$

TABLE 2. Count data

where n is the total number of observations.

Let $D = (r_1, r_2, r_3, r_4)$ be the observed count data in Table 2. Let $\boldsymbol{\theta} = (\theta_1, \theta_2, \theta_3, \theta_4)$ be the vector of the unknown probabilities in Table 1. Then the likelihood function of the unknown probabilities is assumed to follow multinomial distribution:

$$L(\boldsymbol{\theta}) = p(D|\boldsymbol{\theta}) = \frac{(r_1 + r_2 + r_3 + r_4)!}{r_1!r_2!r_3!r_4!} \theta_1^{r_1} \theta_2^{r_2} \theta_3^{r_3} \theta_4^{r_4} \quad (6.1)$$

As a prior for the unknown probabilities we assign the uninformative Dirichlet prior

$$p(\boldsymbol{\theta}) \propto \theta_1^{-1} \theta_2^{-1} \theta_3^{-1} \theta_4^{-1} \quad (6.2)$$

which, if marginalized, collapses to the uninformative Beta prior. Combining the likelihood (6.1) and prior(6.2), by way of the product rule, and normalizing, by way

of the sum rule, we obtain the multivariate Dirichlet posterior distribution of the theta's given the observed count data:

$$p(\boldsymbol{\theta} | D) = \frac{(r_1 + r_2 + r_3 + r_4 - 1)!}{(r_1 - 1)! (r_2 - 1)! (r_3 - 1)! (r_4 - 1)!} \theta_1^{r_1-1} \theta_2^{r_2-1} \theta_3^{r_3-1} \theta_4^{r_4-1} \quad (6.3)$$

Each realization of the function $u(\boldsymbol{\theta})$ maps onto a corresponding probability $p(\boldsymbol{\theta} | D)$. By arranging the values $u(\boldsymbol{\theta})$ on the x -axis and the corresponding $p(\boldsymbol{\theta} | D) d\boldsymbol{\theta}$ on the y -axis, we may obtain the univariate probability distribution of the function $u(\boldsymbol{\theta})$. This probability distribution of $u(\boldsymbol{\theta})$ takes the uncertainty into account we have in regards to the unknown $\boldsymbol{\theta}$ and, consequently, lets us put confidence bounds on this function.

If we only have four unknown theta's we may use brute computational force to partition the domain of $\boldsymbol{\theta}$ and compute of for each partitioning the corresponding pair $[u(\boldsymbol{\theta}), p(\boldsymbol{\theta} | D) d\boldsymbol{\theta}]$, after which we then order the $u(\boldsymbol{\theta})$ and plot them together with their corresponding probabilities $p(\boldsymbol{\theta} | D) d\boldsymbol{\theta}$. However, for large distributions, having many unknown $\boldsymbol{\theta}$'s, this quickly becomes unpractical because of the curse of dimensionality.

In the next section we present an implementation of the Nested Sampling framework by Skilling that allows us to evaluate the probability distribution of $u(\boldsymbol{\theta})$ for highly variate distributions of $\boldsymbol{\theta}$, by way of Monte Carlo sampling. The implementation is accomplished by way of the Inner Nested Sampling algorithm. The Inner Nested Sampling algorithm allows us to sample uniformly from the constrained likelihood space of the $\boldsymbol{\theta}$'s, a necessary prerequisite of the Nested Sampling framework.

7. INNER NESTED SAMPLING

Let L be a likelihood function defined on a highly multivariate parameter space $\boldsymbol{\theta}$. Then Nested Sampling is a Monte Carlo framework with which this multivariate likelihood function $L(\boldsymbol{\theta})$ may be evaluated. The Nested Sampling framework needs uniform samples within the multivariate geometry of some likelihood constraint L^* in order to work. However, this framework does not tell us how to obtain these samples, that is, its optimal implementation is an open-ended research question. In this paper we give an algorithm, called Inner Nested Sampling, that obtains such uniform samples.

The idea behind Inner Nested Sampling is that we obtain a set of differentials of the multivariate geometry of the initial likelihood constraint L^* at iteration step $t = 0$ of Nested Sampling proper. These differentials are defined by a direction \mathbf{e} and a radius $R(\mathbf{e})$ and serve as a proxy for the actual geometry and have the nice property that they may be uniformly sampled. Furthermore as with each iteration step t the geometry defined by likelihood constraint will shrink, the radii $R(\mathbf{e})$

may be updated so as to reflect this shrinkage. This then allows us to continue the uniform sampling of these differentials and, by proxy, the likelihood geometry of interest.

7.1. A change of variable. Say, we have a m -variate parameter vector $\boldsymbol{\theta} = (\theta_1, \dots, \theta_m)$. Then the likelihood constraint $L(\boldsymbol{\theta}) \leq L^*$ defines some sub-domain V^* of the total parameter space V of $\boldsymbol{\theta}$.

Now, if we have an (approximate) modus $\hat{\boldsymbol{\theta}}$ of the likelihood function $L(\boldsymbol{\theta})$ we may translate the origin of $\boldsymbol{\theta}$ to the location of this modus and proceed to make a change of variables from a Cartesian coordinate system to unit-vector coordinates,

$$\boldsymbol{\theta} = \hat{\boldsymbol{\theta}} + r\mathbf{e} \quad (7.1)$$

where $\mathbf{e} = (e_1, \dots, e_m)$ is a point on the unit-sphere and r is the distance from the $\hat{\boldsymbol{\theta}}$ to the likelihood constraint L^* in the direction of \mathbf{e} . Unit-vector coordinates map the parameter vector $\boldsymbol{\theta}$ to a radius r and $(m-1)$ non-redundant coordinates:

$$(\theta_1, \dots, \theta_k) \mapsto (r, e_1, \dots, e_{m-1}) \quad (7.2)$$

where the m th redundant coordinate of the unit vector \mathbf{e} may be found through the identity

$$e_m = 1 - \sum_{i=1}^{m-1} e_i^2$$

The differential of the unit-vector transformation is

$$dV^* = \frac{R(\mathbf{e})^m}{m} dS \quad (7.3)$$

where dS are equi-volume ‘patches’ of the surface S of the m -dimensional unit-sphere. Note that the surface S itself has dimensionality $(m-1)$. The differentials (7.3) are the volumes of pyramids with base dS and height $R(e_1 \cdots e_m)$. So, the integral

$$V^* = \int dV^* = \int_S \frac{R(\mathbf{e})^m}{m} dS \quad (7.4)$$



For example, say, we have the likelihood

$$L(x, y) = 0.184 \exp \left[-\frac{1}{2} (x^2 + xy + y^2) \right] \quad (7.5)$$


which is given in Figure 6. The likelihood constraint $L^* = 0.041$ corresponds with an ellips having ‘volume’ $V^* = 8.162$, Figure 7.

Since we have a two-dimensional likelihood space, the unit sphere is actually an unit circle and it follows that the surface S is actually a circumference, where $S = 2\pi$. The dS is obtained by partitioning the circumference S in n equi-distant line elements, that is $dS = 2\pi/n$. Now, if we let $n = 16$, then

$$dS = \frac{\pi}{8} \quad (7.6)$$

FIGURE 6. Plot of Likelihood L FIGURE 7. Area Defined by Likelihood Constraint $L^* = 0.041$

and the radii R_i , for $i = 1, \dots, n$ at the centers of these line elements dS are shown in Figure 8.

FIGURE 8. Constrained Likelihood Space Partitioned by dS into dV^*

By summing the approximate differentials, (7.3) and (7.6),

$$dV_i^* = \frac{R_i^m}{m} dS = \frac{R_i^m}{m} \frac{\pi}{8} \quad (7.7)$$

we may approximate the integral (7.4):

$$\sum_{i=1}^{16} dV_i^* = 8.151 \approx 8.162 = V^* = \int dV^* \quad (7.8)$$

As we let $n \rightarrow \infty$, this approximation will become evermore accurate.

We summarize, the constrained likelihood space V^* may be represented as a collection of differentials $\{dV_1^*, \dots, dV_n^*\}$. As it will turn out, it is trivially simple to sample uniformly from the set of the differentials dV_i^* , and by doing so we actually sample uniformly the space V^* itself as we let $n \rightarrow \infty$, (7.4).

7.2. Obtaining differentials. The differentials depicted in Figure 8 were obtained through brute computational force. For highly multivariate likelihood spaces such methods are bound to fail due to the curse of dimensionality. However, Skilling's Nested Sampling framework comes here to the rescue, as we shall now demonstrate.

In analogy to the treatment of Nested Sampling proper in Chapter 6, we partition dS in $n = 400$ equi-distant line elements, compute the area of the corresponding differentials $(dV_1^*, \dots, dV_{400}^*)$, and plot them, Figure 9.




FIGURE 9. Area Elements of V^*

Again we are free to reorder these area elements as we like, Figure 10.




FIGURE 10. Ordered Area Elements of V^*

Now all these rectangular elements have a base of $dw = dS = 2\pi/400$. Being that there are 400 area elements we again might view Figure 11 as a representation of some monotonic descending function $g(w)$, where $0 \leq w \leq 2\pi$, Figure 11.




FIGURE 11. Plot of Radii Function g

Were we to integrate this function $g(w)$, we obtain an approximation of the area V^* , (7.8). However, this approximation is not our primary interest in Inner Nested Sampling. Though it may serve a check for those rare cases where we actually know the value of V^* . Rather, the collection of differentials is what we are after, since these may be sampled uniformly, thus, allowing us, by proxy, to uniformly sample V^* itself.

The Nested Sampling framework, previously discussed, translates directly to Inner Nested Sampling, where we set W to be the surface S of the m -dimensional unit sphere

$$W = \frac{m\pi^{m/2}}{\Gamma\left(\frac{m}{2} + 1\right)} \quad (7.9)$$

and take as the function to be evaluated

$$f(\mathbf{e}) = \frac{R(\mathbf{e})^m}{m} = g(w) \quad (7.10)$$

In doing so, we end up with a collection of directions, $\mathbf{e}^{(q)}$, for $q = 1, \dots, Q$, with associated radii $R(\mathbf{e}_q)$, which are the distances from the modus $\hat{\boldsymbol{\theta}}$ to the constraint L^* , and weights $A^{(q)}$, (4.29), (4.32), and (7.10):

$$\begin{aligned} A^{(q)} &= g^{(q)} dw^{(q)} \\ &= g^{(q)} \left(w^{(q-1)} - w^{(q)} \right) \\ &= \frac{R(\mathbf{e}^{(q)})^m}{m} \left[\left(1 - \frac{1}{N+1} \right)^{s-1} - \left(1 - \frac{1}{N+1} \right)^q \right] W \end{aligned} \quad (7.11)$$

where N is the number of objects used in the Inner Nested Sampling run and the N directions \mathbf{e} have been sampled under the constraint

$$R(\mathbf{e}_n^{(q)}) \leq \min R(\mathbf{e}_r^{(q-1)}) \quad n, r = 1, \dots, N \quad (7.12)$$

The weights $A^{(q)}$ correspond with the volumes dV_q^* . By substituting (7.9) in (7.11), we find, (7.3):

$$dV_q^* = A^{(q)} = \frac{R(\mathbf{e}^{(q)})^m}{m} dS_q \quad (7.13)$$

where

$$dS_q = \frac{1}{N+1} \left(1 - \frac{1}{N+1}\right)^{q-1} \frac{m\pi^{m/2}}{\Gamma(\frac{m}{2} + 1)} \quad (7.14)$$

Note that for the differentials dV^* and dS we have let the subscripts enumerate the iteration step.

8. OBTAINING UNIFORM SAMPLES FOR INNER NESTED SAMPLING

The vector $\boldsymbol{\theta} = (\theta_1, \dots, \theta_m)$ is a point in the parameter space and $L(\boldsymbol{\theta})$ is the likelihood function that is defined on this parameter space. In highly dimensional problems the region of interest of $L(\boldsymbol{\theta})$ becomes exponentially small relative to the total sample space. Therefore, the change that random sampling will yield a non-negligible likelihood value becomes exponentially small as the dimensionality of problem grows.

In Inner Nested Sampling not the likelihood function $L(\boldsymbol{\theta})$ is evaluated but the radius function $R(\mathbf{e})$, where \mathbf{e} is a point on the m -dimensional unit sphere. The radius function $R(\mathbf{e})$ is more uniformly spread out over the surface of the unit-sphere than $L(\boldsymbol{\theta})$ is over the parameter space. This makes random sampling of the radius function much more feasible. Nonetheless, as the bound R^* on $R(\mathbf{e})$ becomes tighter and tighter, the number of rejections of the random draws will grow so much as to make a random walk necessary.

8.0.1. *Random sampling of the unit sphere.* Random points on the unit-sphere are easily obtained by letting

$$\mathbf{e} = \frac{\mathbf{e}_0}{\|\mathbf{e}_0\|}, \quad \text{where } \mathbf{e}_0 \sim N(\mathbf{0}, \mathbf{I}_m) \quad (8.1)$$

A rejection occurs whenever $R(\mathbf{e}) < R^*$. If the number of rejections of the random draws of the \mathbf{e} exceeds some number M , where $M \gg 1$, a random walk on the unit sphere is performed.

8.0.2. *Random walk on the unit sphere.* The random walk on the unit sphere takes at its starting point some \mathbf{e}_k that has a radius R_k which is known to satisfy the

radius constraint $R_k > R^*$. and we set

$$\mathbf{e}^{(0)} = \mathbf{e}_k \quad (8.2)$$

If the parameter space is 3-dimensional, then $\mathbf{e}^{(0)}$ will be a point on the surface of a 3-dimensional sphere. Then, for a given length s of the random step to be taken, there is a 2-dimensional circle lying on the surface of this sphere with equi-probable possible candidates $\mathbf{e}^{(1)}$. To obtain a realisation $\mathbf{e}^{(1)}$, we sample a random point \mathbf{u} on this circle, where

$$\mathbf{v} = \frac{\mathbf{v}_0}{\|\mathbf{v}_0\|}, \quad \text{where } \mathbf{v}_0 \sim N(\mathbf{0}, \mathbf{I}_2) \quad (8.3)$$

The circumference of the 2-dimensional circle with the candidates $\mathbf{e}^{(1)}$ will change as the length s of the random step changes. To be more more precise, let $\mathbf{e}^{(1)}$ be the proposal point of the first random step and let α be the angle between the starting point $\mathbf{e}^{(0)}$ and $\mathbf{e}^{(1)}$. Then the radius of the 2-dimensional circle will equal $\sin \alpha$, where $0 \leq \alpha \leq \pi$.

If we translate the origin of our original coordinate system to $(\cos \alpha) \mathbf{e}^{(0)}$ and then rotate the axes of this 3-dimensional system so that the 2-dimensional circle, on which all the $\mathbf{e}^{(1)}$ are to be found, lies in the plane spanned by axes 2 and 3. Then the coordinates of $\mathbf{e}^{(1)}$ in this new coordinate system become, (8.3):

$$\tilde{\mathbf{e}}^{(1)} = \begin{pmatrix} 0 \\ (\sin \alpha) v_1 \\ (\sin \alpha) v_2 \end{pmatrix} \quad (8.4)$$

In order to transform the coordinates of this alternative coordinate system, $\tilde{\mathbf{e}}^{(1)}$, back into the coordinates of the original coordinate system, $\tilde{\mathbf{e}}^{(1)}$, we must realize that in our original coordinate system the 2-dimensional circle of the $\tilde{\mathbf{e}}^{(1)}$ is orthogonal to $(\cos \alpha) \mathbf{e}^{(0)}$. This implies that the rotation matrix B , which accomplishes this re-transformation, has as its first column:

$$\mathbf{b}_1 = \mathbf{e}^{(0)} \quad (8.5)$$

The other two orthonormal columns of B , \mathbf{b}_2 and \mathbf{b}_3 , then can be easily constructed by applying a Gram-Schmidt process. Let

$$Q = \mathbf{I}_3 \quad (8.6)$$

Then

$$\mathbf{b}_2^{(0)} = \mathbf{q}_2 - \frac{\langle \mathbf{b}_1, \mathbf{q}_2 \rangle}{\langle \mathbf{q}_2, \mathbf{q}_2 \rangle} \mathbf{b}_1 = \mathbf{q}_2 - \langle \mathbf{b}_1, \mathbf{q}_2 \rangle \mathbf{b}_1 \quad (8.7)$$

and

$$\mathbf{b}_2 = \frac{\mathbf{b}_2^{(0)}}{\|\mathbf{b}_2^{(0)}\|} \quad (8.8)$$

Likewise, we have

$$\mathbf{b}_3^{(0)} = \mathbf{q}_3 - \langle \mathbf{b}_1, \mathbf{q}_3 \rangle \mathbf{b}_1 - \langle \mathbf{b}_2, \mathbf{q}_3 \rangle \mathbf{b}_2 \quad (8.9)$$

and

$$\mathbf{b}_3 = \frac{\mathbf{b}_3^{(0)}}{\|\mathbf{b}_3^{(0)}\|} \quad (8.10)$$

This results in the rotation matrix

$$B = \begin{pmatrix} \mathbf{e}^{(1)} & \mathbf{b}_2 & \mathbf{b}_3 \end{pmatrix}, \quad \text{where } B^T B = \mathbf{I}_3 \quad (8.11)$$

We then have that

$$\mathbf{e}^{(1)} = (\cos \alpha) \mathbf{e}^{(0)} + B \tilde{\mathbf{e}}^{(1)} \quad (8.12)$$

In the random walk the proposed step $\mathbf{e}^{(t)}$ is rejected whenever either, (7.1),

$$L \left[\hat{\boldsymbol{\theta}} + R \left(\mathbf{e}^{(1)} \right) \mathbf{e}^{(1)} \right] < L^* \quad (8.13)$$

or

$$R \left(\mathbf{e}^{(1)} \right) < R^* \quad (8.14)$$

The random walk step must be modulated in order that the angle α becomes smaller as the number of rejections increases and, visa versa, larger as the number acceptances increases. Now, since a step length of s corresponds with some angle α , we are free to modulate α instead of s . So, let

$$u^{(\tau)} \sim U(0, 1) \quad (8.15)$$

where $\tau \geq t$ and $u^{(\tau)}$ should not equal 0 or 1. Let

$$step^{(\tau)} = \begin{cases} step^{(\tau-1)} + 1/2, & \text{if rejection} \\ step^{(\tau-1)} - 1/2, & \text{if acceptance} \end{cases} \quad (8.16)$$

where $step^{(0)} = 0$. Then, constraining $0 < \alpha^{(\tau)} < \frac{\pi}{2}$, we let

$$\alpha^{(\tau)} = \arcsin \left[\left(u^{(\tau)} \right)^{step^{(\tau)}} \right] \quad (8.17)$$

where it is understood that $u^{(\tau)}$, (8.13), should not equal 0 or 1.

For the case where the dimensionality of the parameter space is $m = 3$, this procedure generalizes to

$$\mathbf{e}^{(t)} = \left(\cos \alpha^{(\tau)} \right) \mathbf{e}^{(t-1)} + B^{(t-1)} \tilde{\mathbf{e}}^{(\tau)} \quad (8.18)$$

where t is the number of random steps accepted in the random walk, τ is the total number of iterations performed by the random walk, and, (8.4),

$$\tilde{\mathbf{e}}^{(\tau)} = \begin{pmatrix} 0 \\ (\sin \alpha^{(\tau)}) v_1^{(\tau)} \\ (\sin \alpha^{(\tau)}) v_2^{(\tau)} \end{pmatrix} \quad (8.19)$$

where, (8.3),

$$\begin{pmatrix} v_1^{(\tau)} \\ v_2^{(\tau)} \end{pmatrix} = \frac{\mathbf{v}_0^{(\tau)}}{\|\mathbf{v}_0^{(\tau)}\|}, \quad \mathbf{v}_0^{(\tau)} \sim N(\mathbf{0}, \mathbf{I}_2) \quad (8.20)$$

In our preliminary empirical studies the number of random walk acceptances fluctuated around the 67%, which is as it should be. Note that this modulation of a random walk on the unit sphere is an adaptation of the modulating algorithm on the Cartesian parameter space as given by Skilling, [1].

8.1. Drawing Uniform samples from Inner Nested Sampling realizations.

If for some m -variate parameter space we have a collection of sampled unit-vectors $\{\mathbf{e}^{(1)}, \dots, \mathbf{e}^{(Q)}\}$ with corresponding radii $\{R^{(1)}, \dots, R^{(Q)}\}$. Then this allows us to draw a uniform sample from the corresponding differentials $\{dV_1, \dots, dV_Q\}$, and by doing so we actually sample uniformly the space V^* itself as we let $Q \rightarrow \infty$, (7.4).

Step 1

First we uniformly draw a differential dV_q . We do this by drawing from the uniform distribution:

$$u \sim U(0, Z^{(Q)}) \quad (8.21)$$

where Z is understood to be the approximation of the V^* , (4.33) and (7.11):

$$Z = \sum_{q=1}^Q A^{(q)} \quad (8.22)$$

We then find the smallest index value q for which the following inequality holds, (7.8):

$$\sum_{i=1}^q A^{(i)} \geq u \quad (8.23)$$

This q then is the index number of the uniformly drawn dV_q .

Step 2

Let

$$v = u - \sum_{i=1}^{q-1} A^{(i)} \quad (8.24)$$

where $0 \leq v \leq A^{(q)}$. Since the differentials dV_q are pyramids in the limit $n \rightarrow \infty$, Figure 8, the realisation v , $0 \leq v \leq dV_q$, geometrically corresponds with the volume of a sub-pyramid having height ρ , where $0 \leq \rho \leq R_q$, that is

$$v = \frac{\rho^m}{m} dS_q \quad (8.25)$$

where (7.14):

$$dS_q = \frac{m\pi^{m/2}}{\Gamma(\frac{m}{2} + 1)} \frac{m}{m+1} \left(1 - \frac{1}{m+1}\right)^{q-1}$$

From (8.25), we then have

$$\rho = \sqrt[m]{\frac{mv}{dS_q}} \quad (8.26)$$

Step 3 The proposal of the uniformly sampled constrained likelihood space of θ , then simply becomes, (7.1):

$$\theta_{\text{proposal}} = \hat{\theta} + \rho \mathbf{e}^{(q)} \quad (8.27)$$

where $\hat{\theta}$ is the modus of the likelihood function which is evaluated, ρ is (8.26), and $\mathbf{e}^{(q)}$ is the direction for which the radius $R^{(q)}$ was determined.

Integrals are limit cases of simple summation. In calculus differentials go to zero in order to obtain infinite precision. So, integrals in general have an infinite amount of differentials. Nested Sampling steps away from the limit case of calculus by approximating an infinite amount of differentials that go to zero with a finite amount of differentials greater than zero, which it then summates. Now, where we to sample to uniformly from V^* then, because of identity (7.4), we may just as well sample uniformly over the infinite amount of differentials dV^* .

Inner Nested Sampling, just like Nested Sampling, also steps away this limit case of infinite precision. By analogy, the finite amount of differentials that were obtained through random sampling are uniformly sampled, rather than the actual infinite amount of differentials. The finite set of differentials, obtained through random sampling, serves as a proxy for the infinite set of differential. So, by sampling uniformly over the finite set of differentials one approximately samples over V^* . The larger the set of differentials, the more accurate this approximation.

8.2. Implementing Nested Sampling. We now give the algorithm that integrates the Inner Nested Sampling within the Nested Sampling framework.

Inner Nested Sampling

Let the Inner Nested Sampling do a run for some large likelihood constraint that is found by shooting a large number of points in some preset parameter space of θ . The smallest likelihood value is taken as the initial likelihood constraint L_0^* .

Let the Inner Nested Sampling algorithm explore the geometry defined by this L_0^* , using M objects. Then we have an approximation of the corresponding domain V^* , in the form of the weights $A^{(q)}$, (7.13), and the associated directions $\mathbf{e}^{(q)}$ and corresponding radii $R^{(q)} = R(\mathbf{e}^{(q)})$.

Nested Sampling proper

Using the procedure in Section 8.1, draw $(N - 1)$ random objects with likelihoods greater than L_0^* from the collection of Inner Nested Sampling differentials $\{dV_1^*, \dots, dV_Q^*\}$ and add these objects with the object corresponding with the initial likelihood constraint L_0^* . We then have N objects with which to perform Nested Sampling proper, as described in Chapter 6.

As the likelihood contour shrinks with every further Nested Sampling iteration, the probability grows that the proposed $\boldsymbol{\theta}$, (8.27), will be rejected. When the ratio of rejections becomes prohibitively large we update the monotonic function g of the Inner Nested Sampling run, Figure 11. This can be done by either determining the new $R^{(q)}$ corresponding with each stored direction vector $\mathbf{e}^{(q)}$. Or, by taking advantage of the fact that the likelihood contour generally will retain its shape as it shrinks, one alternatively may, if pressed for time, take with intervals values of $R^{(q)}$ on the w -‘axis’ and use some interpolating scheme for the other $R^{(q)}$ values. And sampling for Nested Sampling proper then continues for this updated function of g of the Inner Nested Sampling run.

Note that the number of possible directions in the Nested Sampling algorithm is constrained to the number of vectors stored in the Inner Nested Sampling run \mathbf{e} . In order to have many possible directions \mathbf{e} , a large number of objects may be taken in the Inner Nested Sampling run. In our preliminary study we found that for a 20-variate normal likelihood $M = 1000$ initial Inner Nested Sampling objects resulted in $Q = 12.000$ stored direction vectors \mathbf{e} , from which $T = 2000$ random points were actually sampled in the Nested Sampling proper, which was initialized with $N = 100$ objects.

8.3. A caveat. Before we finish the outline of Inner Nested Sampling algorithm we have one more point to make. In Section 8.1, we sampled uniformly from the differential dV_q^* by sampling the corresponding radius $R^{(q)}$. This correspond geometrically with sampling the center line of the differential dV_q^* . Being that this center line itself is sampled uniformly from the unit sphere, either by way of random draw or a random walk, this will still constitute a random draw in case this differential dV_q^* is only sampled once during Nested Sampling proper. However, in the case were we draw the same differential twice, as we sample uniformly over the

collection of differentials, uniformity is compromised, in that given a previous draw over that differential, we have pertinent knowledge of the possible realisations of the samples points.

This may be remedied as follows. At the preliminary phase of Inner Nested Sampling algorithm run, as the radius constraint is not yet so strong that a random walk has to be employed, a very great amount of rejected directions e are generated through pure random sampling. Now, instead of throwing these rejections away, we can store a predesignated number q of the rejections which have radii closest to the radii $R^{(q)}$ already in place in a ‘safety matrix’. This safety matrix is updated as the number of rejections, thus, getting radii which are evermore closer the $R^{(q)}$ already in place.

Now, it is expected that this procedure will go a long way to remedy the possible distortion of uniformity in the sampling within the differentials themselves, since it is especially these differentials dV_q^* in the preliminary phase of Inner Nested Sampling which are most likely to be sampled more than once. As can be deduced from (4.34):

$$dV_q^* = A^{(q)} \rightarrow 0 \quad (8.28)$$

as $q \rightarrow \infty$. However, should it happen that one of the differentials dV_q^* obtained through a random walk is sampled more than once, however unlikely, two choices are left to obtain extra samples for this differentials. Either try to obtain a $\tilde{R}^{(q)}$ by performing a random walk under the constraint

$$\frac{1}{2}R^{(q-1)} \leq \tilde{R}^{(q)} \leq \frac{1}{2}R^{(q+1)} \quad (8.29)$$

whenever in the uniform sampling of the differentials a $R^{(q)}$ is drawn which already has been drawn once. Or, alternatively, ignore possible distortion in uniformity, especially since the dV_q^* more and more become like true differentials, that is, infinitely thin slices of volume, as $q \rightarrow \infty$, (8.28). And the ever mounting computational cost in exploring the radius constraint (8.29) will eventually not outweigh the added benefits that the object $\tilde{R}^{(q)}$ has to offer in terms of the restoration of uniformity, since the distortions in uniformity are so minute.

9. RESOLVING THE ISSUE OF CONFIDENCE BOUNDS

Let $\theta_{ij} = (\theta_{11}, \theta_{12}, \dots, \theta_{IJ})$ for $i = 1, \dots, I$, $j = 1, \dots, J$, where $I \times J = M$. Also let $\theta_{i+} = (\theta_{1+}, \dots, \theta_{I+})$ and $\theta_{+j} = (\theta_{+1}, \dots, \theta_{+J})$. The issue of interest were will apply the Inner Nested Sampling is to evaluate the value $u(\theta_{ij})$. This evaluation is over all plausible values of the θ_{ij} , weighted according to the multivariate posterior distribution of these theta’s, (6.3):

$$p(\theta_{ij} | D) \propto \theta_{11}^{r_{11}-1} \theta_{12}^{r_{12}-1} \dots \theta_{IJ}^{r_{IJ}-1} \quad (9.1)$$

where $D = (r_{11}, r_{12}, \dots, r_{IJ})$.

In the case of a large M the brute computational evaluation of the integrals which give us, respectively, the first and second moment of the function $u(\theta_{ij})$:

$$M_1 = \int p(\theta_{ij}|D) u(\theta_{ij}) d\theta_{ij} \quad (9.2)$$

and

$$M_2 = \int p(\theta_{ij}|D) [u(\theta_{ij})]^2 d\theta_{ij} \quad (9.3)$$

quickly becomes unpractical because of the curse of dimensionality. So, the Nested Sampling framework, as implemented by the Inner Nested Sampling algorithm, will be used to evaluate (9.2). This is done by evaluating, respectively, the sums, (8.19):

$$M_1 \approx \sum_{t=1}^T p(\theta_t) u(\theta_t) = \sum_{t=1}^T \frac{A_t}{Z^{(T)}} u(\theta_t) \quad (9.4)$$

and

$$M_2 \approx \sum_{t=1}^T p(\theta_t) [u(\theta_t)]^2 = \sum_{t=1}^T \frac{A_t}{Z^{(T)}} [u(\theta_t)]^2 \quad (9.5)$$

By way (9.4) and (9.5), we may put the following approximate confidence bounds on the relevance u :

$$M_1 - \sqrt{M_2 - M_1^2} \leq u \leq M_1 + \sqrt{M_2 - M_1^2} \quad (9.6)$$

9.1. Inner Nested Sampling for the Dirichlet Distribution. As it turns out, the dirichlet distribution, that is, the multinomial beta distribution, (9.1) is particularly amenable to the Inner Nested Sampling algorithm, since it has known modulus as well as analytical derivable $R(\mathbf{e})$ for the first run, when no constraints radius constraint is yet in place, other than the simplex form of the domain on which (9.1) is defined. This is particularly helpful, since in general we have that the finding of the radii $R(\mathbf{e})$ is the most computational intensive part of the Inner Nested Sampling, as it requires some kind of search algorithm, such as slice sampling, for example.

We now revert back to the notation of Section 6, and let

$$\theta = \theta_{jk} = \theta_i = (\theta_1, \dots, \theta_M) \quad (9.7)$$

and

$$p(\theta|D) \propto \theta_1^{r_1-1} \dots \theta_M^{r_M-1} \quad (9.8)$$

The modulus of (9.8), using some mathematical package, may be derived as:

$$\hat{\theta} = \hat{\theta}_i = \left(\frac{r_1}{n}, \dots, \frac{r_M}{n} \right) \quad (9.9)$$

where $n = \sum_i r_i$. The radius constraint, for any modulus $\hat{\boldsymbol{\theta}}$, is the distance from this modulus to the boundary of the $(M-1)$ -simplex defined by the vertices

$$\begin{aligned} \mathbf{v}_1 &= (1, 0, \dots, 0) \\ \mathbf{v}_2 &= (0, 1, \dots, 0) \\ &\vdots \\ \mathbf{v}_M &= (0, 0, \dots, 1) \end{aligned} \tag{9.10}$$

where the $(M-1)$ -simplex is the M -dimensional domain of the Dirichlet parameters $\boldsymbol{\theta}$.

The $(M-1)$ -simplex is spanned by the $(M-1)$ vectors, (9.10),

$$\begin{aligned} \mathbf{u}_1 &= \mathbf{v}_1 - \mathbf{v}_M = (1, 0, \dots, 0, -1) \\ \mathbf{u}_2 &= \mathbf{v}_2 - \mathbf{v}_M = (0, 1, \dots, 0, -1) \\ &\vdots \\ \mathbf{u}_{M-1} &= \mathbf{v}_{M-1} - \mathbf{v}_M = (0, 0, \dots, 1, -1) \end{aligned} \tag{9.11}$$

The coordinate system defined by (9.11) is not orthonormal. And in order to make it so, we first orthogonalize the system by way of a Gramm-Schmidt process. First, arbitrarily, let

$$\mathbf{b}_1 = \mathbf{u}_1 \tag{9.12}$$

Then

$$\mathbf{b}_2 = \mathbf{u}_2 - \frac{\langle \mathbf{b}_1, \mathbf{u}_2 \rangle}{\langle \mathbf{u}_2, \mathbf{u}_2 \rangle} \mathbf{b}_1 \tag{9.13}$$

and

$$\mathbf{b}_3 = \mathbf{u}_3 - \frac{\langle \mathbf{b}_1, \mathbf{u}_3 \rangle}{\langle \mathbf{u}_3, \mathbf{u}_3 \rangle} \mathbf{b}_1 - \frac{\langle \mathbf{b}_2, \mathbf{u}_3 \rangle}{\langle \mathbf{u}_3, \mathbf{u}_3 \rangle} \mathbf{b}_2 \tag{9.14}$$

etc., until \mathbf{b}_{M-1} . Then, normalizing the orthogonal vectors (9.14),

$$\mathbf{w}_1 = \frac{\mathbf{b}_1}{\|\mathbf{b}_1\|} \quad \dots \quad \mathbf{w}_{M-1} = \frac{\mathbf{b}_{M-1}}{\|\mathbf{b}_{M-1}\|} \tag{9.15}$$

we obtain the $(M-1)$ -dimensional orthonormal coordinate system W we are looking for:

$$\mathbf{W} = \begin{bmatrix} \mathbf{w}_1 & \dots & \mathbf{w}_{M-1} \end{bmatrix} \tag{9.16}$$

Let \mathbf{e} be a random point on the $M-1$ -dimensional unit-sphere, obtained by letting, (8.1),

$$\mathbf{e} = \frac{\mathbf{e}_0}{\|\mathbf{e}_0\|}, \quad \text{where } \mathbf{e}_0 \sim N(\mathbf{0}, \mathbf{I}_{M-1}) \tag{9.17}$$

And let

$$\boldsymbol{\delta} = \mathbf{W} \mathbf{e} \tag{9.18}$$

Then any point

$$\mathbf{p} = \hat{\boldsymbol{\theta}} + R \boldsymbol{\delta}, \quad R \geq 0 \quad (9.19)$$

will lie on the same ‘plane’ as the $(M-1)$ -simplex. In order to get the maximum radius R for which \mathbf{p} is a point on the the boundary of the $(M-1)$ -simplex, we must realize that on this boundary one of the elements of the \mathbf{p} must necessarily be zero. So we solve for every element of \mathbf{p} :

$$p_i = \hat{\boldsymbol{\theta}}_i + R_i \boldsymbol{\delta}_i = 0 \quad (9.20)$$

and obtain

$$R_i = -\frac{\hat{\boldsymbol{\theta}}_i}{\boldsymbol{\delta}_i} \quad (9.21)$$

As only positive radii R_i are allowed, the maximum radius R for which \mathbf{p} still lies on the boundary of the $(M-1)$ -simplex is

$$R = \min R_i, \quad \text{for those } R_i \geq 0 \quad (9.22)$$

As any radius greater than R will automatically result in negative p_i values, which then would imply that \mathbf{p} lies outside of the $(M-1)$ -simplex.

We summarize, in order to use Inner Nested Sampling on the M -dimensional Dirichlet distribution (9.8), we draw a random point on the $(M-1)$ -dimensional unit-sphere, (9.17), which we then transform to the direction (9.18). And the corresponding distance R from the modus $\hat{\boldsymbol{\theta}}$ to the boundary of the domain of the Dirichlet parameters, $\boldsymbol{\theta}$, which is the $(M-1)$ -simplex defined by the vertices (9.10), is found by way of (9.21) and (9.22). Furthermore, the volume of this $(M-1)$ -simplex is known to be

$$V^* = \frac{\sqrt{M}}{(M-1)!} \quad (9.23)$$

and (9.23) then may serve as check on the accuracy of the estimation, (8.19),

$$\log Z^{(T)} \approx \log V^* \approx (M-1) \log (M-1) - (M-1) - \frac{1}{2} \log M \quad (9.24)$$

REFERENCES

- [1] Skilling, J.: Nested Sampling, In Bayesian Inference and Maximum Entropy Methods in Science and Engineering, (eds. Erickson G., Rychert J.T., and Smith C.R.) AIP Conference Proceedings, American Institute of Physics, New-York. 395-405, (2004).
- [2] Skilling, J.: Nested Sampling for General Bayesian Computation, Bayesian Analysis 4, pp. 833-860, (2006).

- [3] Skilling, J.: The Canvas of Rationality, In Bayesian Inference and Maximum Entropy Methods in Science and Engineering, Sao Paulo, Brazil 2008 (eds. Lauretto MS, Pereira CAB) AIP Conference Proceedings, American Institute of Physics, New-York. 67-79, (2008).
- [4] Zellner A.: An Introduction to Bayesian Inference in Econometrics, J. Wiley & Sons, Inc., New York, (1971).
- [5] MacKay D.J.C.: Information Theory, Inference, and Learning Algorithms, Cambridge University Press, Cambridge, (2003).